

The Arms Race: Geopolitics of the Buildout

Shay O'Kelly · August 2026

The last essay ended on an observation that deserves its own essay. When a handful of physical chokepoints decide which nations get to manufacture intelligence and how fast, those chokepoints stop being industrial assets and become strategic ones. The twentieth century organized its geopolitics around oil, and we fought wars, built alliances, and drew maps around the places it came out of the ground. Intelligence is now on the same path, except the concentration is more extreme than oil ever was. There was never a moment when one company in one country made every barrel, but there is exactly one company on earth that makes the lithography machines behind every advanced chip, one foundry that manufactures nearly all of them, and it sits on an island claimed by the country America is racing. When national growth rates depend on a supply chain this concentrated, foreign policy becomes supply chain management, and that is not a metaphor. It is a literal description of what export controls are.

On paper, the chokepoint map is a Western royal flush. ASML is Dutch. TSMC is Taiwanese and fabs mostly for American designers. The frontier labs are American, the hyperscaler capital is American, and the high-bandwidth memory is Korean. Starting in 2022, Washington began playing that hand through export controls designed to deny China advanced chips and the tools to make them, and the logic followed directly from the scaling laws in the first essay. If capability is a function of compute, and you control compute, you control capability. It is the most aggressive use of economic chokepoints since the Cold War, and it rests on one assumption, that the denied input cannot be substituted or engineered around.

The evidence since then says the assumption is only half true, and the half that failed is instructive. In January 2025, DeepSeek released a reasoning model trained under embargo, on restricted hardware, that matched the leading Western models at a fraction of the training cost, and Nvidia lost almost six hundred billion dollars of market value in a single day processing the news, the largest one-day loss in market history¹. Since then the pattern has repeated on a schedule. Moonshot's Kimi K3, released this July at 2.8 trillion parameters, is the largest open-source model in the world². Eight of the top ten Chinese models are open-weight, and Alibaba's Qwen family has passed Meta's Llama in cumulative downloads to become the default open model for much of the world's developers³. The embargo was supposed to ration capability, but capability is a function of compute and ideas, and rationing one raises the return on the other. Constraint became a forcing function. Chinese labs got world-class at efficiency precisely because they were not allowed to be world-class at scale.

And notice what China does with those models, it gives them away. This is not idealism. It is the logic of the last essay weaponized. If you cannot monopolize the frontier, commoditize it, because every workflow in the world that runs on a free Chinese open-weight model is a workflow that generates no revenue for the American labs trying to earn back a trillion dollars of capex, and a workflow that runs on weights whose

training data and assumptions were shaped in Beijing. America is export-controlling atoms while China is export-subsidizing bits. One of those strategies gets stronger as models commoditize, and it is not ours.

The hardware story is more honest for the West, but the direction of travel matters more than the snapshot. Huawei's Ascend chips, the flagship of Chinese self-sufficiency, turned out on teardown to be built substantially on TSMC dies acquired through sanctions evasion, nearly three million of them routed through a shell buyer⁴. Chinese self-sufficiency in 2026 is partly theater. But look at what is happening underneath the theater. A state-backed group in Shanghai has begun low-volume production of domestic immersion DUV lithography machines, with around five shipping this year to SMIC, Hua Hong, and CXMT, and twenty planned next year. Against ASML's 131 immersion tools shipped last year⁵, that is nothing, and the Chinese machines are worse, slower, and still dependent on imported Japanese components. It is also the hardest step in the entire supply chain being taken for the first time. DUV with multi-patterning prints chips a generation or two behind the frontier, expensively, and expensive-but-domestic is exactly the kind of problem command economies are built to brute-force. On a decade view, the safe assumption is not that the moat holds. It is that the moat buys time, and the question becomes what each side does with the time.

Which brings me to the number that I think matters more than any model release, and that almost nobody in the AI debate talks about. In 2025, China added 543 gigawatts of new power capacity, an energy buildout of roughly half a trillion dollars in a single year⁶. The entire American grid added about 60. That is roughly nine to one against the country running the most power-hungry buildout in history, and China's additions since 2021 alone exceed the entire installed US grid. Meanwhile American data centers added 8.5 gigawatts of capacity last year, have 13.6 scheduled this year⁷, and demand is projected to more than double from 31 gigawatts to 66 by 2027⁸, straight into multi-year transformer lead times, sold-out turbine slots, and interconnection queues that move at the speed of litigation. Now put this next to the bottleneck migration from the last essay. The constraint on manufactured intelligence is moving from silicon, where America holds the chokepoints, toward power, where China does. Export controls are a fortress built around the bottleneck of 2024, while the bottleneck of 2030 is being cornered by the other side. If the physical constraint keeps migrating toward energy, time is not on our side, and the most important industrial policy question in America is not chip subsidies. It is why one gigawatt takes five years to connect.

Having made the case for worry, I should be clear about where I actually land, because it is not on the doom side. I am more bullish on the United States, and the reason is the substitutability rule from the last essay applied to nations instead of stocks. America's deficit is energy, and energy has workarounds. Gas behind the meter, solar and storage in the desert, siting compute where stranded power already exists, restarting nuclear plants. And the workaround list keeps growing. Elon Musk says he wants solar collected in orbit within the decade, where the sun never sets and nobody files an environmental review, and whether or not his timeline holds, the direction is the point. The menu of ways around the energy constraint gets longer every year. Closing an energy gap is an engineering and permitting problem, ugly but solvable with money and will, and America has both, plus some of the cheapest natural gas on earth. China's deficit is the semiconductor supply chain, and that has no workaround. To close it, China has to recreate ASML, TSMC,

and the hundreds of specialized suppliers beneath them, decades of accumulated process knowledge across the most intense supply chain humans have ever built, starting from five domestic DUV machines a year against 131. Both powers are racing to fix their weakness, but the weaknesses are not symmetric. America is short a commodity. China is short a miracle. That is why I believe the US semiconductor lead outweighs China's energy lead, and why the real threat to the American position is not China's grid. It is our own permitting.

Underneath the scoreboard is a more uncomfortable question. Which operating system is better suited to a buildout like this, authoritarianism or democracy? The honest answer cuts both ways. Command systems are built for pouring concrete. When the state decides power gets built, permits do not take five years, transmission lines do not die in court, and coal, solar, nuclear, and hydro get built simultaneously at whatever scale the plan demands. Democracies are built for allocating capital, and it shows. The American buildout is the largest privately financed infrastructure project in history, funded by price signals rather than decrees, and every layer of the frontier stack was invented inside open societies. But each system is now being tested at its weak point. Authoritarian buildouts misallocate on a colossal scale and answer to no market when they are wrong. Democratic buildouts underbuild and litigate, and the veto points that make democracies humane in normal times, the environmental review, the local objection, the appeal, have quietly become the binding constraint on democratic AI capacity. The race will partly be decided by which failure mode costs more this decade. And the stakes are not symmetric, because manufactured intelligence is not a neutral input to both systems. Cheap, abundant cognition is the most powerful surveillance and social-control technology ever created, which makes it a regime-strengthening technology for authoritarian states in a way it is not for open ones. The values embedded in the models the world runs on, and the system that builds the most capacity to run them, are the same question viewed from two angles.

Follow the state money and the securitization is already visible on every continent. Washington has moved from subsidy toward ownership, with \$52.7 billion of CHIPS grants, an equity stake in Intel itself, a \$500 billion public-private compute buildout in Stargate, and a Department of Energy contract to restart domestic uranium enrichment after the country outsourced essentially all of that capacity at the end of the Cold War⁹. Beijing runs the same play, a \$47.5 billion third national semiconductor fund stacked on state compute funds and a state-dominated rare earth chain. Brussels has a €43 billion Chips Act and a €200 billion AI investment program. Japan is deploying roughly \$65 billion, Korea tens of billions more, and the Gulf is building gigawatt AI campuses with sovereign wealth¹⁰. Seven of the largest American technology companies have taken nearly \$38 billion of government support over their lifetimes, SpaceX and Tesla among them⁹. Every layer of the stack, in every bloc, now has state capital behind it. That is what an arms race looks like on a balance sheet.

So where does it land? The last essay argued that on a decade view models commoditize toward the cost of tokens, with one exception. If AI becomes a first-order geopolitical concern, commoditization stops at national borders. I now think that exception is the base case at the frontier. The race dynamics push both governments the same direction. As capability starts to look militarily decisive, frontier models get treated less like software and more like enriched uranium, walled inside national security perimeters, while the commodity layer, the Kimis and Qwens and open-weight Llamas, diffuses globally and prices toward zero.

Two stacks, two blocs, a commodity floor and a classified frontier. And here is the investment observation hiding in the geopolitics. Every one of those futures, commoditized or walled, cooperative or cold-war, requires more physical capacity, because a securitized race builds harder than a commercial one. Nations do not capex-discipline their way through an arms race. The bottleneck thesis is not a bet on which bloc wins. It is a bet that both keep building, and that is the safest bet in the whole series.

One consequence of an arms race in manufactured intelligence stays hidden in the national accounting, and it is what the race does to the people inside each country. The same buildout that redistributes power between nations redistributes income within them, away from wages and toward the owners of the machines, at a speed no tax system was designed for. That is the subject of the next essay.

SOURCES

1. DeepSeek R1 release (January 20, 2025) and Nvidia's \$589B single-day market-cap loss (January 27, 2025). Via Forbes and MIT Technology Review.
2. Kimi K3 (2.8 trillion parameters, July 2026). Moonshot AI release coverage and AI Proem.
3. Chinese open-weight share (8 of top 10) and Qwen passing Llama in downloads. Via Inference Hub, MIT Technology Review, and HuggingFace download data.
4. Huawei Ascend teardowns showing TSMC dies and the ~2.9M die Sophgo purchase. TechInsights via SemiconductorX.
5. China domestic immersion DUV production (~5 tools 2026, ~20 planned 2027, deliveries to SMIC, Hua Hong, CXMT) and ASML's 131 immersion tools shipped. Reuters reporting via Tom's Hardware and TrendForce.
6. China power additions (543 GW in 2025, ~\$500B buildout). China National Energy Administration figures via OilPrice.com and CarbonCredits.com.
7. US data center capacity additions (8.5 GW realized 2025, 13.6 GW scheduled 2026) via Goldman Sachs Research (Aterio data). US grid additions (~60 GW in 2025) and China's post-2021 additions exceeding the installed US grid via EIA and Bloomberg, January 2026.
8. US data center power demand (31 GW in 2025 to 66 GW by 2027). Goldman Sachs Research.
9. US state backing. CHIPS Act grants (\$52.7B program), the US government equity stake in Intel (Intel 8-K, August 2025), Stargate (\$500B public-private compute buildout), the DOE 10-year ~\$900M HALEU enrichment contract (General Matter, Paducah), and ~\$38B lifetime subsidies across the seven largest US tech companies incl. SpaceX and Tesla. Via the US Department of Commerce, Intel SEC filings, the US Department of Energy, and press reports.
10. China Big Fund III (\$47.5B state semiconductor fund) and rare-earth supply chain control via Caixin and SCMP. EU Chips Act (€43B) and InvestAI (€200B) via the European Commission. Japan ~\$65B chip programs incl. Rapidus (~\$18B) via METI and Bloomberg. Korea semiconductor support (\$19B+) via Chosun Daily. Gulf, the HUMAIN 1GW cluster within the 5GW US-UAE Stargate campus, via Arab News and state announcements.

No Offer or Solicitation. The only purpose of this site is to provide general background information on RAXIA. This is not an offer to sell or the solicitation of any offer to buy any securities. Any such offer would be made only to qualified investors after distribution and return of the Private Placement Memorandum and governing agreement for the applicable Fund (the "Offering Materials"). Nothing on

this site is investment advice or a recommendation of any security. The author may hold positions in companies discussed. Past performance is not indicative of future results.